The High-Performance Storage System

Staking a Position at the Leading Edge

HPSS

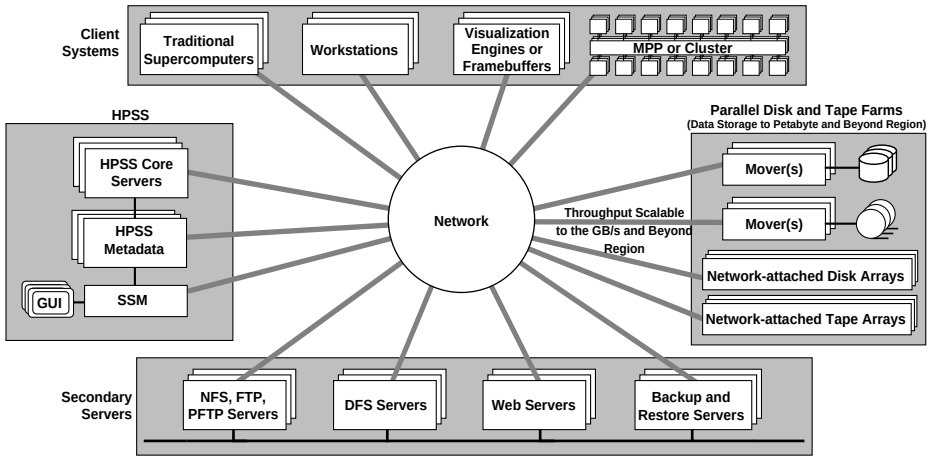

Example of the system architecture supported by HPSS.

Technology

The rapid growth in the size of datasets has caused a serious imbalance in I/O and storage system performance and functionality relative to application requirements and the capabilities of other system components. The High-Performance Storage System (HPSS) is a scalable, next-generation storage system that is meeting the functionality and performance requirements of large-scale scientific and commercial computing environments.

Impact

Our goal is to improve the performance and capacity of hierarchical storage systems by two orders of magnitude or more over what is available in the general marketplace today. We also provide corresponding improvements in architecture and functionality.

The rapidly increasing requirements in computational science, enterprise integration, multimedia, and data collection plus rapid improvements in processing capability, main memory sizes, and networking are producing very large datasets. These datasets range from tens of gigabytes up to terabytes. In the near future, they will require total capacities scalable into the petabyte range and beyond, both distributed and at single sites. We expect these large datasets and capacities to be common in high-performance and large-scale scientific and commercial environments within the national information infrastructure. Results of this rapid growth of data are a serious imbalance in I/O and storage system performance, functionality relative to application requirements, and the capabilities of other system components.

To deal with these issues, the HPSS collaboration is working to improve the performance and capacity of storage systems by two orders of magnitude or more over what is available in the general marketplace today, with corresponding improvements in architecture and functionality.

The HPSS Collaboration

TheHPSS collaboration is based on the premise that no single organization has the experience and resources to meet all the challenges represented by the growing storage system I/O, capacity, and functionality imbalance described above. Over 20 organizations including industry, Department of Energy (DOE) and other federal laboratories, universities, and National Science Foundation (NSF) supercomputer centers have contributed to various aspects of this effort. The HPSS development team consists of IBM Government Systems, four DOE laboratories (Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia), and NASA's Langley Research Center. Transarc is collaborating on integrating HPSS and the Distributed File System (DFS). Cornell University and NASA's Lewis Research Center also helped with software development.

In addition to the above DOE laboratories and NASA Langley, other early deployment partners include the California Institute of Technology and the Jet Propulsion Laboratory, Cornell Theory Center, Fermi National Accelerator Laboratory, Maui High Performance Computer Center, San Diego Supercomputer Center, and the University of Washington. Each of the early deployment partners have challenging heterogeneous environments and research and development agendas that add significant value to HPSS. Argonne National Laboratory, National Center for Atmospheric Research, and Pacific Northwest Laboratory have contributed requirements and other assistance to the work.

Industry hardware and software contributors to the collaboration have included Ampex, IBM, Storage Technology, Kinesix, Network Systems Corp., PsiTech, Sony Precision Graphics, and Zitel.

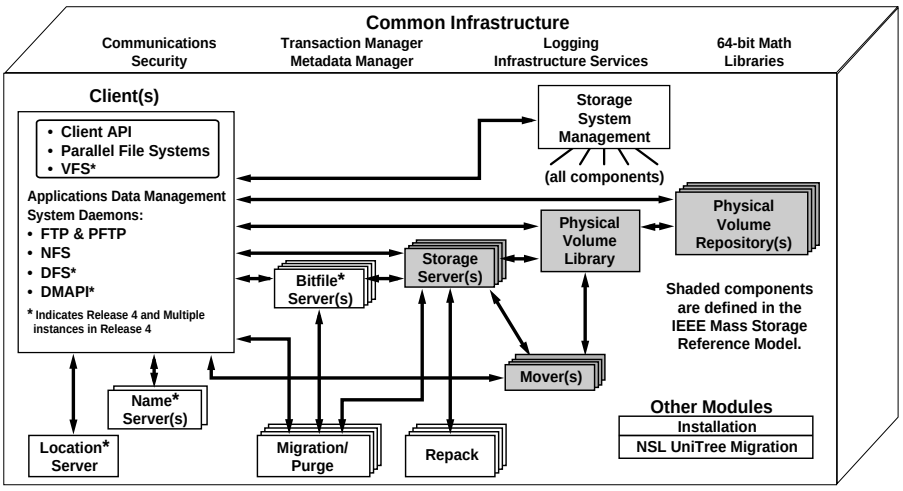

HPSS software architecture. To improve modularity, we are building on the IEEE Mass Storage Reference Model architecture.

The HPSS Objectives

The HPSS is designed to support configurations such as the one shown above. Several requirements are driving the HPSS design and implementation:

- Scalability in several dimensions—for example, distribution and multiprocessing of servers, data transfer rates to gigabytes per second and beyond through support for network-attached peripherals and parallel I/O, storage capacity to petabytes and beyond, unlimited file sizes , number of naming directories to millions, and hundreds to thousands of simultaneous clients.
- Modularity by building on the IEEE Reference Model architecture (see the figure above), to support client access to all major system subcomponents, replacement of software components during the storage system's lifecycle, and integration of multi-vendor hardware and software storage components.
- Portability to many vendors' platforms by building on industry standards, such as the Open Software Foundation's (OSF's) Distributed Computing Environment (DCE), standard communications protocols, C, POSIX, and UNIX with no kernel modifications. HPSS uses commercial products for system infrastructure, and all HPSS module interfaces have been given to the IEEE Storage System Standards Working Group.
- Reliability and recoverability through support for atomic transactions among distributed components, mirroring and logging of system metadata or user data, recovery from failed devices or media, reconnection logic, ability to relocate distributed components, and use of software engineering development practices.
- Client APIs to all major system components, including a parallel client API and interface to vendor parallel file systems, and support for industry standard services such as FTP (sequential and parallel) and NFS. Support will also include DFS and DMAPI (DMIG).
- Security through DCE and POSIX security mechanisms, including authentication, access control lists, file permissions, and security labels.
- System manageability through a managed object reporting, monitoring, and database framework; and management operations with graphical user interface access and control.
- Support for better integration with data management and other application systems through appropriate interface functionality at multiple levels in the architecture.
- Distributability by building on a client/server architecture and use of an OSF DCE infrastructure.

HPSS Commercialization

The HPSS commercialization plan includes availability and support by IBM as a high-end Solution Service offering through IBM Government Systems. HPSS source code may also be licensed and marketed by any U.S. company.

For more information about HPSS, contact Dick Watson, LLNL, 510-422-9216, dwatson@llnl.gov; or Bob Coyne, IBM Government Systems, 713-335-4040, coyne@vnet.ibm.com.